

Model Checking Classes of Metric LTL Properties of Object-Oriented Real-Time Maude Specifications

Daniela Lepri
University of Oslo, Norway

Peter Csaba Ölveczky
University of Oslo, Norway

Erika Ábrahám
RWTH Aachen University, Germany

This paper presents a transformational approach for model checking two important classes of metric temporal logic (MTL) properties, namely, bounded response and minimum separation, for non-hierarchical object-oriented Real-Time Maude specifications. We prove the correctness of our model checking algorithms, which terminate under reasonable non-Zeno-ness assumptions when the reachable state space is finite. These new model checking features have been integrated into Real-Time Maude, and are used to analyze a network of medical devices and a 4-way traffic intersection system.

1 Introduction

Real-Time Maude [20] is a formal specification language and a high-performance simulation and model checking tool that extends the rewriting-logic-based Maude system [8] to support the formal specification and analysis of *real-time* systems. Real-Time Maude differs from timed-automaton-based tools, such as UPPAAL [5] and KRONOS [27], by emphasizing ease and expressiveness of specification over algorithmic decidability of key properties. In particular, Real-Time Maude supports the definition of any computable data type, unbounded data structures, different communication models, and so on.

Because of its expressiveness, Real-Time Maude has been successfully applied to a wide range of advanced state-of-the-art applications that are beyond the pale of timed automata, including the OGDC density control [22] and LMST topology control [10] protocols for wireless sensor networks, the CASH scheduling algorithm with capacity sharing features that require unbounded queues [16], the AER/NCA active networks multicast protocol [21], and the NORM multicast protocol developed by the IETF [13]. Real-Time Maude's natural model of time, together with its expressiveness, also makes it ideal as a semantic framework in which real-time modeling languages can be given a formal semantics; such languages then also get Real-Time Maude's formal analysis capabilities essentially for free. Languages with a Real-Time Maude semantics include: a timed extension of the Actor model [9], the Orc web services orchestration language [2], a language developed at DoCoMo laboratories for handset applications [1], a behavioral subset of the avionics standard AADL [15], the visual model transformation language e-Motions [25], real-time model transformations in MOMENT2 [6], and a subset of Ptolemy II discrete-event models [4].

Real-Time Maude is particularly suitable to model real-time systems in an *object-oriented* style, and the paper [20] identifies some useful specification techniques for object-oriented real-time systems. All the concrete applications mentioned above, and many of the language semantics applications, are specified in an object-oriented way using those techniques.

Real-Time Maude provides a spectrum of analysis methods, including simulation through timed rewriting, untimed temporal logic model checking, and (unbounded or time-bounded) search for reachability analysis. However, up to now, Real-Time Maude has lacked the ability to model check *timed* (or *metric*) temporal logic properties. Such properties are obviously very important in many real-time

systems. For example, in case of an accident the airbag must not just inflate *eventually*, but within very tight time bounds. For timed automata, such metric temporal logic model checking is decidable¹, and implemented in the KRONOS tool [27]. For the much more expressive Real-Time Maude formalism, supporting metric temporal logic checking, is obviously a much harder task.

This paper reports on our first attempts at providing metric temporal logic model checking for Real-Time Maude. We have taken the following pragmatic choices:

1. Supporting the model checking of only a few classes of metric temporal logic properties, namely, the ones that were needed in the above-mentioned applications. These properties are:
 - *Bounded response*: each p -state must be followed by a q -state within time r (where p and q are state propositions). One example of a bounded response property is “whenever the ventilator assisting the patient’s breathing is turned off, it must be turned on within 5 seconds”.
 - *Minimum separation*: there must be at least time r between two non-consecutive p -states. For example, “the ventilator should be turned on continuously for at least two minutes between two pauses.”
2. Supporting such model checking only for flat object-oriented models specified according to the guidelines mentioned above. But as already said, this class of systems includes all the concrete Real-Time Maude applications listed above.

What is gained by restricting the classes of systems and properties is *efficiency*. Instead of implementing the model checking algorithms from scratch, we pursue a *transformational* approach, where we take advantage of Maude’s high performance analysis commands and transform a metric model checking problem $\mathcal{R}, L, t_o \models \phi$ into a problem $\tilde{\mathcal{R}}, \tilde{L}, \tilde{t}_o \models \tilde{\phi}$ that can be analyzed by Real-Time Maude’s efficient search and LTL model checking commands. Our transformations add a clock which measures, respectively, the time since the earliest $(p \wedge \neg q)$ -state that has not been followed by a q -state (for bounded response) and the last time since we saw a p -state (for minimum separation). An important property is that – under reasonable time-divergence assumptions about the executions with the selected time sampling strategy – if the original reachable state space is finite, then the model checking commands are guaranteed to terminate. Furthermore, our model checking commands are semi-decision procedures for the invalidity of the metric properties for time-diverging systems. The transformations have been implemented in Real-Time Maude and the corresponding model checking commands have been made available in the tool. We have applied the new commands on two case studies, one on the safe interoperation of medical devices [14] and one on a fault-tolerant controller for traffic lights in an intersection [17].

We prove the correctness of the transformation under reasonable assumptions, such as the real-time rewrite theory being *tick-invariant* [19]. Since real-time rewrite theories do not have a “region-automaton”-like discrete quotient, for dense time Real-Time Maude uses *time sampling strategies* to execute the tick rules. That is, in model checking analyses for dense-time models, only a subset of all possible behaviors are analyzed. Therefore, Real-Time Maude analyses are in general not (both) sound and complete; however, for object-oriented specifications we have identified easily checkable conditions that guarantee soundness and completeness of our analyses also for dense-time systems [19].

This paper is organized as follows. Section 2 introduces Real-Time Maude and metric temporal logic. Section 3 presents the properties that we address and the corresponding transformations, whose correctness is proved in Section 4. Section 5 shows two case studies of metric temporal logic model checking in Real-Time Maude. Section 6 discusses related work, and Section 7 gives some concluding remarks.

¹for finite behaviour, see, e.g., [7]

2 Preliminaries

2.1 Real-Time Maude

In Real-Time Maude [20], real-time systems are modeled by a set of *equations* and *rewrite rules*. The rewrite rules are divided into *instantaneous* rules, that model changes that are assumed to take zero time, and *tick* rules that model time advance. Formally, a Real-Time Maude *timed module* specifies a *real-time rewrite theory* [18] of the form $\mathcal{R} = (\Sigma, E, IR, TR)$, where:

- (Σ, E) is a *membership equational logic* [8] theory with Σ a signature² and E a set of *confluent and terminating conditional equations*. (Σ, E) specifies the system's state space as an algebraic data type, and must contain a specification of a sort `Time` modeling the (discrete or dense) time domain. We denote by $\mathbb{T}_{\mathcal{R}, s}$ all ground terms of sort `s`.
- IR is a set of (possibly conditional) *labeled instantaneous (rewrite) rules* specifying the system's *instantaneous* (i.e., zero-time) local transitions, written `cr1 [l] : t => t' if cond`, where l is a *label*. Such a rule specifies a *one-step transition* from an instance of t to the corresponding instance of t' . The rules are applied *modulo* the equations E .³
- TR is a set of *tick (rewrite) rules*, written with syntax
`cr1 [l] : {t} => {t'} in time τ if cond .`
 that model time elapse. $\{ _ \}$ is a built-in constructor of sort `GlobalSystem`, and τ is a term of sort `Time` that denotes the *duration* of the rewrite.

The initial state must be a ground term of sort `GlobalSystem` and must be reducible to a term of the form $\{t\}$ using the equations in the specification. The form of the tick rules ensures that time advances uniformly in the whole system.

Following [18], we write $t \xrightarrow{r} t'$ when t can be rewritten into t' in time r by a *one-step rewrite*. Note that instantaneous steps have duration 0. A (*timed*) *path* π in \mathcal{R} is an infinite sequence

$$\pi = t_0 \xrightarrow{r_0} t_1 \xrightarrow{r_1} t_2 \dots$$

such that either

- for all $i \in \mathbb{N}$, $t_i \xrightarrow{r_i} t_{i+1}$ is a one-step rewrite in \mathcal{R} ; or
- there exists a $k \in \mathbb{N}$ such that $t_i \xrightarrow{r_i} t_{i+1}$ is a one-step rewrite in \mathcal{R} for all $0 \leq i < k$, there is no one-step rewrite from t_k in \mathcal{R} , and $t_j = t_k$ and $r_{j-1} = 0$ for each $j > k$.

We denote by $Paths(\mathcal{R})_{t_0}$ the set of all timed paths of \mathcal{R} starting in t_0 . We call a path $\pi = t_0 \xrightarrow{r_0} t_1 \xrightarrow{r_1} t_2 \dots$ *time-divergent* iff for all $r \in \mathbb{R}$ there is an $i \in \mathbb{N}$ such that $\sum_{k=0}^i r_k > r$. Paths that are not time-divergent are called *time-convergent*. We define $\pi^k = t_k \xrightarrow{r_k} t_{k+1} \xrightarrow{r_{k+1}} \dots$. A term t' is *reachable* from t_0 in \mathcal{R} in time r iff there is a path $\pi = t_0 \xrightarrow{r_0} \dots \xrightarrow{r_{k-1}} t_k \dots$ with $t_k = t'$ and $r = \sum_{i=0}^{k-1} r_i$.

The Real-Time Maude syntax is fairly intuitive; we refer to [8] for a detailed description. For example, a function symbol f is declared with the syntax `op f : s1 ... sn -> s`, where $s_1 \dots s_n$ are the sorts of its arguments, and s is its (value) *sort*. Equations are written with syntax `eq t = t'`, and `ceq t = t' if cond` are conditional equations. The mathematical variables in such statements are declared with the keywords `var` and `vars`.

In *object-oriented* Real-Time Maude modules, a *class* declaration

²That is, Σ is a set of declarations of *sorts*, *subsorts*, and *function symbols*.

³ E is a union $E' \cup A$, where A is a set of equational axioms such as associativity, commutativity, and identity, so that deduction is performed *modulo* A . Operationally, a term is reduced to its E' -normal form modulo A before any rewrite rule is applied.

```
class C | att1 : s1, ... , attn : sn .
```

declares a class C with attributes att_1 to att_n of sorts s_1 to s_n , respectively. An *object* of class C in a state is represented as a term $\langle O : C \mid att_1 : val_1, \dots, att_n : val_n \rangle$ of sort `Object`, where O , of sort `Objid`, is the object's *identifier*, and where val_1 to val_n are the current values of the attributes att_1 to att_n , respectively. In a *concurrent* object-oriented system, the state is a term of sort `Configuration`. It has the structure of a *multiset* made up of objects and messages. Multiset union for configurations is denoted by a juxtaposition operator (empty syntax) that is declared associative and commutative, so that rewriting is *multiset rewriting* supported directly in Real-Time Maude. The dynamic behavior of concurrent object systems is axiomatized by specifying its transition patterns by rewrite rules. For example, the rule

```
r1 [l] : m(0,w) < 0 : C | a1 : 0, a2 : y, a3 : w > =>
      < 0 : C | a1 : T, a2 : y, a3 : y + w > dly(m'(0'),x) .
```

defines a parametrized family of transitions (one for each substitution instance), which can be applied whenever the attribute a_1 of an object 0 of class C has the value 0 , with the effect of altering the attributes a_1 and a_3 of the object. Moreover, a message m , with parameters 0 and w , is read and consumed, and a new message $m'(0')$ is sent *with delay* x (see [20]). “Irrelevant” attributes, such as a_2 , need not be mentioned in a rule.

A *flat* (or *non-hierarchical*) object-oriented specification is one where all rewrites happen in the “outermost” configuration; that is, no attribute value t rewrites to some $t' \neq t$.

The specification of time-dependent behavior of object-oriented real-time systems follows the techniques given in [20]. Time elapse is modeled by the tick rule

```
var C : Configuration .   var T : Time .
cr1 [tick] : {C} => {delta(C, T)} in time T if T <= mte(C) [nonexec] .
```

The function `delta` defines the effect of time elapse on a configuration, and the function `mte` defines the maximum amount of time that can elapse before some action must take place. These functions distribute over the objects and messages in a configuration and must be defined for all single objects and messages to define the timed behavior of a system. The tick rule advances time *nondeterministically* by *any* amount T less than or equal to `mte(C)`. To execute such rules, Real-Time Maude offers a choice of *time sampling strategies*, so that only *some* moments in time are visited. The choice of such strategies includes:

- Advancing time by a fixed amount Δ in each application of a tick rule.
- The *maximal* strategy, that advances time to the next moment when some action must be taken, as defined by `mte`. This corresponds to *event-driven simulation*.

Formal Analysis. A Real-Time Maude specification is *executable*, under reasonable conditions, and the tool offers a variety of formal analysis methods. The *rewrite* command simulates *one* fair behavior of the system *up to a certain duration*. The *search* command uses a breadth-first strategy to analyze all possible behaviors of the system, by checking whether a state matching a *pattern* and satisfying a *condition* can be reached from the initial state. Such a pattern typically describes the *negation* of an invariant, so that the search succeeds iff the invariant is violated. The command which searches for n states satisfying the *pattern* search criterion has syntax

```
(utsearch [n] t =>* pattern such that cond .)
```

Real-Time Maude also extends Maude's *linear temporal logic model checker* to check whether each behavior, possibly up to a certain time bound, satisfies a temporal logic formula. *State propositions* are terms of sort `Prop`, and their semantics should be given by (possibly conditional) equations of the form

$\{statePattern\} \models prop = b$

for b a term of sort `Bool`, which defines the state proposition $prop$ to hold in all states $\{t\}$ where $\{t\} \models prop$ evaluates to `true`. We use the notation Π for the set of propositions and L_Π for the (implicit) labeling function assigning to each state the set of propositions that hold in the state. A temporal logic *formula* is constructed by state propositions and the Boolean and temporal logic operators discussed in Section 2.2. The time-bounded model checking command has syntax

$(mc\ t \models_t formula\ in\ time\ \leq \tau\ .)$

for initial state t and temporal logic formula $formula$.

Since the model checking commands execute tick rules according to the chosen time sampling strategy, only a subset of all possible behaviors is analyzed. Therefore, Real-Time Maude analyses are in general *incomplete* for a given property. However, in [19] we have given easily checkable conditions for ensuring that Real-Time Maude analyses are indeed sound and complete.

It is also worth remarking that in the rest of the paper, we implicitly consider the different analyses w.r.t. Real-Time Maude executions. That is, for dense time, by “a rewrite theory \mathcal{R} ” in the following sections we typically mean the real-time rewrite theory \mathcal{R}^{tss} that has been obtained from an original time-nondeterministic real-time rewrite theory \mathcal{R} by applying the theory transformation corresponding to using the time sampling strategy tss when executing the tick rules [20].

2.2 Metric Temporal Logic

Linear temporal logic (LTL) [24] allows us to describe properties of paths of a given system. The states are labeled with elements from a finite set Π of atomic propositions. Besides propositions and the usual Boolean operators, LTL formulae can be built using the temporal *until* operator. Intuitively, the formula $p\ U\ q$ (“ p until q ”) is satisfied by a path if the property q becomes valid within an arbitrary but finite number of steps and the property p constantly holds on the path before. As syntactic sugar we define $\Diamond p$ (“eventually p ”, defined as $true\ U\ p$) that is satisfied by a path if p holds somewhere on the path, and $\Box p$ (“globally p ”, defined as $\neg(true\ U\ (\neg p))$) expressing that p holds on the whole path. The *weak until* operator $p\ W\ q$ is defined as $(p\ U\ q) \vee (\Box p)$.

For time-critical systems we need more expressive power to state that some actions should happen *within some time bounds*. There are different extensions of LTL to capture also timed properties (see [3] for an overview). In this paper, we use the extension *metric temporal logic (MTL)* [11], that adds time interval bounds to the temporal operators. For the until operator, the formula $p\ U_{[t_1, t_2]}\ q$ states that $p\ U\ q$ holds and, furthermore, q occurs within the time interval $[t_1, t_2]$.

Formulae of MTL are built using the following abstract syntax:

$$\varphi ::= true \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi\ U_{[t_1, t_2]}\ \varphi$$

with $p \in \Pi$ and either $t_1, t_2 \in \mathbb{R}$ with $t_1 \leq t_2$ and $t_2 > 0$, or $t_1 \in \mathbb{R}$ and $t_2 = \infty$. Note that $U_{[0, \infty]}$, for which we just write U , corresponds to the unbounded until of LTL. Besides the usual Boolean operators \vee, \rightarrow, \dots we define as syntactic sugar $\Diamond_{[t_1, t_2]}\ \varphi$ as $true\ U_{[t_1, t_2]}\ \varphi$ and $\Box_{[t_1, t_2]}\ \varphi$ as $\neg(true\ U_{[t_1, t_2]}\ (\neg\varphi))$. If the lower bound t_1 is 0, we use the notation $\varphi_1\ U_{\leq t_2}\ \varphi_2$, and analogously for the other operators.

Given a real-time rewrite theory \mathcal{R} , the set of *states* is defined as $\mathbb{T}_{\Sigma/E, GlobalSystem}$. A set Π of (possibly parametric) *atomic propositions* on those states can be defined equationally in a protecting extension $(\Sigma \cup \Pi, E \cup D) \supseteq (\Sigma, E)$, and give rise to a *labeling function* $L_\Pi : \mathbb{T}_{\Sigma/E, GlobalSystem} \rightarrow \mathcal{P}(\Pi)$ in the obvious way [8]. Adapting the pointwise semantics for MTL given in [3], we can define satisfaction of MTL formulas for real-time rewrite theories over timed paths as follows:

Definition 2.1. Let \mathcal{R} be a real-time rewrite theory, L_Π a labeling function on \mathcal{R} , and $\pi = t_0 \xrightarrow{r_0} t_1 \xrightarrow{r_1} \dots$ a timed path in \mathcal{R} . The satisfaction relation of an MTL formula ϕ for the path π in \mathcal{R} is then defined recursively as follows:

$$\begin{aligned}
\mathcal{R}, L_\Pi, \pi &\models \text{true} && \text{always holds} \\
\mathcal{R}, L_\Pi, \pi &\models p && \text{iff } p \in L_\Pi(t_0) \\
\mathcal{R}, L_\Pi, \pi &\models \neg \phi && \text{iff } \mathcal{R}, L_\Pi, \pi \not\models \phi \\
\mathcal{R}, L_\Pi, \pi &\models \phi_1 \wedge \phi_2 && \text{iff } \mathcal{R}, L_\Pi, \pi \models \phi_1 \text{ and } \mathcal{R}, L_\Pi, \pi \models \phi_2 \\
\mathcal{R}, L_\Pi, \pi &\models \phi_1 \text{ } U_{[r_a, r_b]} \text{ } \phi_2 && \text{iff there exists a } j \in \mathbb{N} \text{ such that } \mathcal{R}, L_\Pi, \pi^j \models \phi_2, \\
&&& \mathcal{R}, L_\Pi, \pi^i \models \phi_1 \text{ for all } 0 \leq i < j, \text{ and } r_a \leq \sum_{k=0}^{j-1} r_k \leq r_b.
\end{aligned}$$

For a state t_0 of sort `GlobalSystem`, the satisfaction relation of an MTL formula ϕ for the state t_0 in \mathcal{R} is defined as:

$$\mathcal{R}, L_\Pi, t_0 \models \phi \iff \forall \pi \in \text{Paths}(\mathcal{R})_{t_0} \quad \mathcal{R}, L_\Pi, \pi \models \phi$$

3 Model Checking MTL Properties of Object-Oriented Specifications

Real-Time Maude currently does not support MTL model checking. However, some MTL formulas can already be model checked in Real-Time Maude using the *time-bounded* search and LTL model checking commands. For example, we can model check the time-bounded until property $\mathcal{R}, L_\Pi, t_0 \models p \text{ } U_{\leq r} \text{ } q$, for p and q state properties from Π , using the time-bounded model checking command

```
(mc t0 |=t p U q in time <= r .)
```

We can also analyze the properties $\mathcal{R}, L_\Pi, t_0 \models \Box_{\leq r} p$ and $\mathcal{R}, L_\Pi, t_0 \models \Diamond_{\leq r} p$ in a similar way.

In this paper we present analysis algorithms for the following two classes of MTL formulae:

1. *Bounded response:* $\Box (p \rightarrow (\Diamond_{\leq r} q))$
2. *Minimum separation:* $\Box (p \rightarrow (p \text{ } W (\Box_{\leq r} \neg p)))$

We propose to transform an MTL model checking problem $\mathcal{R}, L_\Pi, t_0 \models \phi$ into an untimed LTL model checking problem $\tilde{\mathcal{R}}, \tilde{L}_\Pi, \tilde{t}_0 \models \tilde{\phi}$. Both transformations add a *clock* to the system: for model checking bounded response properties, this clock measures the time since p held without q holding in the meantime; for minimum separation properties, the clock measures the distance between two non-consecutive p -states. We take care not to increase the clocks “unnecessarily,” so that if the state space reachable from t_0 in \mathcal{R} is finite, then the state space reachable from \tilde{t}_0 in $\tilde{\mathcal{R}}$ remains finite, under reasonable time-divergence assumptions on the executions.

We assume that our specifications are *tick-invariant* [19] with regard to the state propositions occurring in the formula, i.e., a tick step does not change the valuation of the atomic propositions occurring in the formula. Most systems, including the two case studies in the paper, satisfy tick-invariance, since the state propositions usually do not involve the value of clock and timer attributes in the system.

3.1 Bounded response: $\Box (p \rightarrow \Diamond_{\leq r} q)$

A bounded response property states that the system always reacts to a request p with an action q within time r . For example, in our medical devices case study, the ventilation machine, helping a sedated patient to breathe, should not be stopped for more than two seconds at a time; that is, each state in which the machine is pausing must be followed by a state in which the machine is breathing in two seconds or less.

The MTL model checking problem

$$\mathcal{R}, L_{\Pi}, t_0 \models \Box (p \longrightarrow \Diamond_{\leq r} q)$$

for $p, q \in \Pi$ state propositions, can be transformed into the untimed model checking problem

$$\tilde{\mathcal{R}}_r, \tilde{L}_{\Pi}, \tilde{t}_0 \models \Box (p \longrightarrow \Diamond q) \wedge \Box (\text{clock}(c_{BR}) \leq r)$$

where $\text{clock}(c_{BR})$ is the value of a “clock” that measures the time since p held without q holding in the meantime. For real-time rewrite theories having only time-divergent paths we could skip the first condition $\Box (p \longrightarrow \Diamond q)$, that assures, that we also consider all relevant time-convergent paths as possible counterexamples.

We add a “clock” c_{BR} to the system, and update it as follows:

- i) If the clock c_{BR} is turned off, and a state satisfying $p \wedge \neg q$ is reached, then the clock is set to 0 and is turned on.
- ii) The clock is turned off when a state satisfying q is reached.
- iii) A clock that is on is increased according to the elapsed time in the system.

For the very useful class of “flat” object-oriented specifications formalized according to the guidelines in [20]—all advanced Real-Time Maude applications have been so specified—we can automate the transformation from $\mathcal{R}, L_{\Pi}, t_0, p, q, r$ to $\tilde{\mathcal{R}}, \tilde{L}_{\Pi}, \tilde{t}_0$ as follows:

1. Add the following class for the clock:

```
class Clock | clock : Time, status : OnOff .
sort OnOff .      ops on off : -> OnOff [ctor] .
```

2. Add a clock object to the initial state $\{t_0\}$, so that the initial state becomes

$$\{t_0 \quad < c_{BR} : \text{Clock} \mid \text{clock} : 0, \text{status} : x >\}$$

where c_{BR} is a constant of sort `0id` and x is `on` if $p \in L(\{t_0\})$ and $q \notin L(\{t_0\})$, and is `off` otherwise. Note that $p \in L(\{t_0\})$ can be checked in Maude by checking whether $\{t_0\} \models p = \text{true}$.

3. We keep Real-Time-Maude’s object-oriented tick rule and extend the functions `delta` and `mte` to clocks as follows, ensuring that `mte` is not affected by the new clock object:

```
eq delta(< c_{BR} : Clock | status : on, clock : T >, T') =
    < c_{BR} : Clock | clock : if T <= r then T + T' else T fi > .
eq delta(< c_{BR} : Clock | status : off >, T') = < c_{BR} : Clock | > .
eq mte(< c_{BR} : Clock | >) = INF .
```

Notice that the `delta` function ensures that the clock value never increases more than necessary, preserving *finiteness* of the reachable state space from the initial state.

4. Each *instantaneous* rule $t \Rightarrow t'$ if *cond* or $\{t\} \Rightarrow \{t'\}$ if *cond* in \mathcal{R} is replaced by the rules:

```
{t REST < c_{BR} : Clock | status : on >}
=> {t' REST < c_{BR} : Clock | >} if {t' REST} |= q /= true and cond
```

(if the clock is on, then it continues to stay on if a state satisfying $\neg q$ is reached);

```
{t REST < cBR : Clock | status : on >}
=> {t' REST < cBR : Clock | status : off >} if {t' REST} |= q and cond
```

(if the clock is on, then it is turned off when a state satisfying q is reached);

```
{t REST < cBR : Clock | status : off >}
=> {t' REST < cBR : Clock | clock : 0, status : on >}
   if {t' REST} |= p and {t' REST} |= q /= true and cond
```

(if the clock is off, then it is set to 0 and turned on when a state satisfying $p \wedge \neg q$ is reached);

```
{t REST < cBR : Clock | status : off >}
=> {t' REST < cBR : Clock | >}
   if {t' REST} |= q or {t' REST} |= p /= true and cond
```

(if the clock is off, then it continues to stay off if a state satisfying $q \vee \neg p$ is reached).

In the above rules REST is a variable of sort Configuration that does not appear in the original rule. REST matches the “other” objects and messages in the state.

Summarizing, the *BR-transformation* transforms a real-time rewrite theory \mathcal{R} , a labeling function L_Π of \mathcal{R} with $p, q \in \Pi$, an initial state t_0 of \mathcal{R} , and a bounded response formula $\Box(p \longrightarrow \Diamond_{\leq r} q)$ into the triplet $\tilde{\mathcal{R}}, \tilde{L}_\Pi$, and \tilde{t}_0 by

- transforming \mathcal{R} into $\tilde{\mathcal{R}}$ according to the points 1, 3, and 4 above;
- transforming L_Π into \tilde{L}_Π by adapting its domain to the transformed state space, but letting the labeling otherwise unchanged, i.e., $L_\Pi(\{t\}) = \tilde{L}_\Pi(\{t\ o\})$ for all states t of \mathcal{R} and all Clock instances o ;
- extending the initial state t_0 according to point 2 above, yielding \tilde{t}_0 .

The validity of the bounded response property $\Box(p \longrightarrow \Diamond_{\leq r} q)$ is equivalent to $\Box(p \longrightarrow \Diamond q)$ and the clock value being less than or equal to r in each reachable state of the transformed module. The latter property can be defined as an atomic proposition

```
op clock'<=_ : Time -> Prop [ctor] .
eq {REST < cBR : Clock | clock : T1 >} |= clock <= T2 = (T1 <= T2) .
```

and hence bounded response can be analyzed using Real-Time Maude’s untimed LTL model checking features. We have implemented the above model transformation in Real-Time Maude. We have also implemented a bounded response model checking command in the tool based on this transformation. However, for pragmatic reasons, we do *not* model check the property $\tilde{\mathcal{R}}, \tilde{L}_\Pi, \tilde{t}_0 \models \Box(p \longrightarrow \Diamond q) \wedge \Box(\text{clock}(c_{BR}) \leq r)$. Instead, we have observed the unsurprising fact that, with time sampling strategy executions, all our large Real-Time Maude applications are modeled as time-diverging theories. In these cases, bounded response reduces to checking $\tilde{\mathcal{R}}, \tilde{L}_\Pi, \tilde{t}_0 \models \Box(\text{clock}(c_{BR}) \leq r)$, which can be analyzed by the following search command that searches for a state in which the clock value is greater than r :

```
(utsearch [1] {t0 < cBR : Clock | clock : 0, status : x >} =>*
   {C:Configuration < cBR : Clock | clock : T:Time >} such that T:Time > r .)
```

where x is on if $p \in L(\{t_0\})$ and $q \notin L(\{t_0\})$, and is off otherwise. The practical difference is that, whereas the LTL model checking does not terminate when the state space reachable from t_0 in \mathcal{R} is infinite, the above search command provides a *semi-decision* procedure for the invalidity of the bounded response property. For an example of the benefit of this time-divergence-assuming implementation,

consider the bounded response analysis of the medical systems example in Section 5. The reachable state space is infinite because of the clock used in the original model; hence any direct LTL model checking would not terminate, but we see that our bounded response command indeed returns a counterexample falsifying the bounded response property.

In our tool, the bounded response model checking command (for the automatic *BR*-transformation and the execution the Real-Time Maude search) is written with syntax

```
(br  $t_0$  | =  $p \Rightarrow \langle \rangle \text{le}(r) q$  .)
```

3.2 Minimum Separation: $\Box (p \rightarrow (p \text{ W } \Box_{\leq r} \neg p))$

Given a real-time rewrite theory \mathcal{R} with a labeling function L_Π , $p \in \Pi$, all runs of \mathcal{R} are made up of a sequence of blocks for which p and $\neg p$ hold alternately (see Figure 1). The minimum separation property requires that each $\neg p$ -block occurring after a p -block must have a minimum duration r . I.e., if the run for which we check the property starts with a p -block, then all $\neg p$ -blocks of the run must have a duration at least r . Otherwise, if the run starts with a $\neg p$ -block, then the same holds for all $\neg p$ -blocks except the first one at the beginning of the run.




Figure 1: The form of runs satisfying the minimum separation property $\Box (p \rightarrow (p \text{ W } \Box_{\leq r} \neg p))$. The p - and $\neg p$ -blocks may also be infinite.

We transform the MTL model checking problem

$$\mathcal{R}, L_\Pi, t_0 \models \Box (p \rightarrow (p \text{ W } \Box_{\leq r} \neg p))$$

into the untimed model checking problem

$$\tilde{\mathcal{R}}, \tilde{L}_\Pi, \tilde{t}_0 \models \Box (\text{status}(c_{MS}) = \text{on} \vee \text{clock}(c_{MS}) \geq r)$$

where $\text{clock}(c_{MS})$ is the value of a “clock” that measures the time duration since we saw a p -state. That means, to model check minimum separation properties, we add a “clock” c_{MS} to the system, which is initially turned off and set to r : in this way we ensure that an eventual initial $\neg p$ -block does not cause a violation of the property. We update the clock as follows:

- i) If we move from a p -state to a $\neg p$ -state, then the clock is turned on and reset to 0.
- ii) The clock is turned off when a state satisfying p is reached.
- iii) A clock that is on is increased according to the elapsed time in the system.

We can automate the transformation to search for counterexamples of a minimum separation property of the above form as follows:

1. Add the same class for the clock as in Section 3.1:

```
class Clock | clock : Time, status : OnOff .
```

2. Add a clock object to the initial state $\{t_0\}$, yielding

```
 $\{t_0 < c_{MS} : \text{Clock} \mid \text{clock} : r, \text{status} : \text{off} >\}$ 
```

where c_{MS} is a constant of sort `Objid`.

3. We keep Real-Time-Maude's object-oriented tick rule and extend the function `delta` and `mte` to clocks exactly as in Section 3.1.

4. Each *instantaneous* rule $t \Rightarrow t'$ if *cond* or $\{t\} \Rightarrow \{t'\}$ if *cond* in \mathcal{R} is replaced by the rules:

```
 $\{t \text{ REST} < c_{MS} : \text{Clock} \mid \text{status} : \text{on} >\}$   

 $\Rightarrow \{t' \text{ REST} < c_{MS} : \text{Clock} \mid >\}$  if  $\{t' \text{ REST}\} \models p \neq \text{true}$  and cond
```

(if the clock is on, then it continues to stay on, if a state satisfying $\neg p$ is reached);

```
 $\{t \text{ REST} < c_{MS} : \text{Clock} \mid \text{status} : \text{on} >\}$   

 $\Rightarrow \{t' \text{ REST} < c_{MS} : \text{Clock} \mid \text{status} : \text{off} >\}$  if  $\{t' \text{ REST}\} \models p$  and cond
```

(if the clock is on, then it is turned off when a state satisfying p is reached);

```
 $\{t \text{ REST} < c_{MS} : \text{Clock} \mid \text{status} : \text{off} >\}$   

 $\Rightarrow \{t' \text{ REST} < c_{MS} : \text{Clock} \mid >\}$   

  if  $(\{t \text{ REST}\} \models p \neq \text{true}$  or  $\{t' \text{ REST}\} \models p)$  and cond
```

(the clock remains off, if either we are in a state satisfying $\neg p$ or we move to a state satisfying p ; the first condition is needed to avoid switching the clock on in initial $\neg p$ -blocks);

```
 $\{t \text{ REST} < c_{MS} : \text{Clock} \mid \text{status} : \text{off} >\}$   

 $\Rightarrow \{t' \text{ REST} < c_{MS} : \text{Clock} \mid \text{status} : \text{on}, \text{clock} : 0 >\}$   

  if  $\{t \text{ REST}\} \models p$  and  $\{t' \text{ REST}\} \models p \neq \text{true}$  and cond
```

(if the clock is off, and we move from a state satisfying p to a state satisfying $\neg p$, then the clock is turned on and reset to 0).

Again, `REST` is a variable of sort `Configuration` that does not appear in the original rule.

The *MS-transformation* therefore transforms a real-time rewrite theory \mathcal{R} , a labeling function L_Π with $p \in \Pi$, an initial state t_0 of \mathcal{R} , a state proposition p , and a time value r into the triple $\tilde{\mathcal{R}}, \tilde{L}_\Pi$, and \tilde{t}_0 by

- transforming \mathcal{R} into $\tilde{\mathcal{R}}$ according to the points 1, 3, and 4 above;
- transforming L_Π into \tilde{L}_Π by adapting its domain to the transformed state space, but letting the labeling otherwise unchanged, i.e., $\tilde{L}_\Pi(\{t \ o\}) = L_\Pi(\{t\})$ for all states t of \mathcal{R} and all `Clocks` o ;
- extending the initial state t_0 according to point 2 above, yielding \tilde{t}_0 .

Checking the minimum separation property $\Box (p \rightarrow (p \ W \ \Box_{\leq r} \neg p))$ is equivalent to checking that the validity of p implies that the clock value is larger than or equal to r in each state in the transformed module. The violation of the latter can be checked by the following search command that searches for a state in which the clock is off (which implies that p holds) and the clock value is smaller than r :

```
(utsearch [1]  $\{t_0 < c_{MS} : \text{Clock} \mid \text{clock} : r, \text{status} : \text{off} >\} \Rightarrow *$   

  {C:Configuration < c_{MS} : Clock | clock : T:Time, status : off >}  

  such that T:Time < r .)
```

The above *MS-transformation* has been integrated in Real-Time Maude, and model checking the above minimum separation property can be done with the Real-Time Maude command

```
(ms  $t_0 \models p$  separated by  $\geq r$  .)
```

4 Correctness of Bounded Response Model Checking

In this section we give the correctness proof for our bounded response model checking. The correctness proof for minimum separation, which we omit due to lack of space, is quite similar, and can be found in an extended version of this paper [12].

To increase readability, in the following we use the notation $\pi \models \phi$ instead of $\mathcal{R}, L_\Pi, \pi \models \phi$ if \mathcal{R} and L_Π are clear from the context.

The following lemma states that the *BR*-transformation only adds some observers to the original systems, without modifying its behavior.

Lemma 4.1. *Let \mathcal{R} be a real-time rewrite theory, L_Π with $p, q \in \Pi$ a labeling function for \mathcal{R} , and let $\{t_0\}$ be an initial state for \mathcal{R} . Let $\tilde{\mathcal{R}}, \tilde{L}_\Pi$, and $\{\tilde{t}_0\}$ be the result of the *BR*-transformation applied to \mathcal{R}, L_Π , and t_0 .*

Then for each path $\{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots$ in \mathcal{R} there is a path $\{\tilde{t}_0\} \xrightarrow{r_0} \{\tilde{t}_1\} \xrightarrow{r_1} \dots$ in $\tilde{\mathcal{R}}$ such that, for all i , there exists t'_i with $\tilde{t}_i = t_i \ t'_i$ and vice versa, for all paths $\{\tilde{t}_0\} \xrightarrow{r_0} \{\tilde{t}_1\} \xrightarrow{r_1} \dots$ in $\tilde{\mathcal{R}}$ there is a path $\{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots$ in \mathcal{R} such that, for all i , $\tilde{t}_i = t_i \ t'_i$ for some t'_i .

Proof. Adding the clock class and a clock object to the initial state does not affect the original part of the state, and defining mte of the additional clocks to be the infinity value INF ensures that the new clocks don't modify the timed behavior of the (original) system. Furthermore, the transformation replaces each original rule by a number of new rules, such that (1) each new rule acts on the original state part as the original rule, and (2) for each original rule and each extended state to which the original rule is applicable there is exactly one new rule that is applicable. (1) assures that the new rewrites yield the same result for the original part of the state and (2) assures that no original paths are blocked by the new rules. Thus the transformation does not modify the original behavior.

“ \rightarrow ”: Let $\{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots$ be a path of \mathcal{R} . We define

$$\tilde{t}_i = t_i < c_{BR} : \text{Clock} \mid \text{clock} : x_i, \text{status} : y_i >$$

for all i with $x_i \in \mathbb{T}_{\mathcal{R}, \text{Time}}$ and $y_i \in \mathbb{T}_{\mathcal{R}, \text{OnOff}}$ given inductively as follows:

- $x_0 = 0$, and $y_0 = \text{on}$ if $p \in L_\Pi(\{t_0\}) \wedge q \notin L_\Pi(\{t_0\})$ and $y_0 = \text{off}$ otherwise.
- For all i , if there is a tick rule yielding the rewrite $\{t_i\} \xrightarrow{r_i} \{t_{i+1}\}$, then we distinguish between the following cases:
 - If $y_i = \text{on}$ and $x_i \leq r$, then we define $y_{i+1} = \text{on}$ and $x_{i+1} = x_i + r_i$.
Note that with the definition of the `delta` equation we have $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$.
 - If $y_i = \text{on}$ and $x_i > r$, then we define $y_{i+1} = \text{on}$ and $x_{i+1} = x_i$.
Note that with the definition of `delta` we have $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$.
 - Else, if $y_i = \text{off}$, we define $y_{i+1} = \text{off}$ and $x_{i+1} = x_i$.
With the definition of the `delta` equation we have $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$.
- For all i , otherwise there is an instantaneous rule $t \Rightarrow t'$ if *cond* or $\{t\} \Rightarrow \{t'\}$ if *cond*, yielding the rewrite $\{t_i\} \xrightarrow{r_i} \{t_{i+1}\}$ with $r_i = 0$.
 - If $y_i = \text{on}$ and $\{t_{i+1}\} \models q \neq \text{true}$ then we set $y_{i+1} = \text{on}$ and $x_{i+1} = x_i$.
Note that the first replacement of the original rule yields $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$.
 - If $y_i = \text{on}$ and $\{t_{i+1}\} \models q$ then we set $y_{i+1} = \text{off}$ and $x_{i+1} = x_i$.
Note that the second replacement of the original rule yields $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$.

- If $y_i = \text{off}$, $\{t_{i+1}\} \models p$, and $\{t_{i+1}\} \not\models q \neq \text{true}$ then we set $y_{i+1} = \text{on}$ and $x_{i+1} = 0$.
Note that the third replacement of the original rule yields $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$.
- Else, if $y_i = \text{off}$ and either $\{t_{i+1}\} \models q$ or $\{t_{i+1}\} \models p \neq \text{true}$ then we set $y_{i+1} = \text{off}$ and $x_{i+1} = x_i$.
Note that the fourth replacement of the original rule yields $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$.

Above we made use of the fact that by definition for each i , the corresponding labeling $\tilde{L}_\Pi(\{\tilde{t}_i\})$ in $\tilde{\mathcal{R}}$ is equal to $L_\Pi(\{t_i\})$. Clearly, all $\{\tilde{t}_i\}$ are states of $\tilde{\mathcal{R}}$. Especially, $\{\tilde{t}_0\}$ results from $\{t_0\}$ by the *BR*-transformation. Thus $\{\tilde{t}_0\} \xrightarrow{r_0} \{\tilde{t}_1\} \xrightarrow{r_1} \dots$ is a path of $\tilde{\mathcal{R}}$.

“ \leftarrow ”: Given a path $\{\tilde{t}_0\} \xrightarrow{r_0} \{\tilde{t}_1\} \xrightarrow{r_1} \dots$ of $\tilde{\mathcal{R}}$ such that

$$\tilde{t}_i = t_i < c_{BR} : \text{Clock} \mid \text{clock} : x_i, \text{status} : y_i >$$

for each i , we show that $\{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots$ is a path of \mathcal{R} .

- For all i , if $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$ can be gained by a tick rule in $\tilde{\mathcal{R}}$, then clearly also $\{t_i\} \xrightarrow{r_i} \{t_{i+1}\}$ can be gained by a tick rule in \mathcal{R} .
- Otherwise if $\{\tilde{t}_i\} \xrightarrow{r_i} \{\tilde{t}_{i+1}\}$ can be gained by an instantaneous rule in $\tilde{\mathcal{R}}$, then the original rule which got replaced by the above one yields $\{t_i\} \xrightarrow{r_i} \{t_{i+1}\}$ in \mathcal{R} .

■

The following lemma clarifies the semantics of the bounded response property: On the one hand, if along a path after a p event r time long no q event occurs, then the path is a counterexample for the property. On the other hand, if a path violates the bounded response property, then either after a p event r time long no q event occurs, or the path is time-convergent and violates the unbounded property $\Box(p \rightarrow (\Diamond q))$.

Lemma 4.2. *Let \mathcal{R} be a real-time rewrite theory, L_Π with $p, q \in \Pi$ a labeling function for \mathcal{R} , and $\pi = \{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots$ a path of \mathcal{R} . Then*

$$\left[\exists i, j. 0 \leq i < j \wedge (\pi^i \models p) \wedge (\forall i \leq k \leq j. \pi^k \not\models q) \wedge \sum_{k=i}^{j-1} r_k > r \right] \longrightarrow [\pi \not\models \Box(p \rightarrow (\Diamond_{\leq r} q))]$$

and

$$[\pi \not\models \Box(p \rightarrow (\Diamond_{\leq r} q))] \longrightarrow \left[\exists i, j. 0 \leq i < j \wedge (\pi^i \models p) \wedge (\forall i \leq k \leq j. \pi^k \not\models q) \wedge \sum_{k=i}^{j-1} r_k > r \right] \vee [\pi \not\models \Box(p \rightarrow (\Diamond q))].$$

Proof. For the first implication, due to the semantics of MTL the following holds:

$$\begin{aligned} & \exists i, j. 0 \leq i < j \wedge (\pi^i \models p) \wedge (\forall i \leq k \leq j. \pi^k \not\models q) \wedge \sum_{k=i}^{j-1} r_k > r. \longrightarrow \\ & \exists i. (\pi^i \models p) \wedge \forall j \geq i. \left(\sum_{k=i}^{j-1} r_k \leq r \rightarrow \pi^j \not\models q \right) \longrightarrow \\ & \exists i. (\pi^i \models p) \wedge (\pi^i \not\models \Diamond_{\leq r} q) \longrightarrow \\ & \exists i. \pi^i \models \neg(p \rightarrow (\Diamond_{\leq r} q)) \longrightarrow \\ & \pi \not\models \Box(p \rightarrow (\Diamond_{\leq r} q)). \end{aligned}$$

For the other direction,

$$\begin{aligned} \pi \not\models \Box (p \rightarrow (\Diamond_{\leq r} q)) &\rightarrow \exists i. \pi^i \models \neg(p \rightarrow (\Diamond_{\leq r} q)) \\ &\rightarrow \exists i. (\pi^i \models p) \wedge (\pi^i \not\models \Diamond_{\leq r} q) \\ &\rightarrow \exists i. (\pi^i \models p) \wedge \forall j \geq i. \left(\sum_{k=i}^{j-1} r_k \leq r \rightarrow \pi^j \not\models q \right). \end{aligned}$$

Let i be such an index with $\pi^i \models p$ and $\forall j \geq i. (\sum_{k=i}^{j-1} r_k \leq r \rightarrow \pi^j \not\models q)$. If $\pi \not\models \Box (p \rightarrow (\Diamond q))$ then we are ready. So assume $\pi \models \Box (p \rightarrow (\Diamond q))$, implying that there is a smallest index $l \geq i$ with $\pi^l \models q$. From the above it follows that $\sum_{k=i}^l r_k > r$.

Note that by definition $r > 0$ and thus $l > i$. Let $j = l - 1$. From the minimality of l we first conclude that $\forall i \leq k \leq j. \pi^k \not\models q$. From the minimality of l we furthermore conclude that the rewrite $\{\tilde{t}_j\} \rightarrow \{\tilde{t}_l\}$ is an instantaneous step, and thus $\sum_{k=i}^j r_k = \sum_{k=i}^l r_k > r$. That means,

$$\exists i, j. 0 \leq i < j \wedge (\pi^i \models p) \wedge \left(\forall i \leq k \leq j. \pi^k \not\models q \right) \wedge \sum_{k=i}^{j-1} r_k > r.$$

■

The following main theorem formalizes the correctness of our transformation: Firstly, if the bounded response property holds, then the model checking algorithms will not provide any counterexample. Secondly, if the bounded response model checking algorithm does not find any counterexample, and if there are no time-convergent counterexamples, then the property holds.

Theorem 4.3. *Let \mathcal{R} be a real-time rewrite theory, L_Π a labeling function for \mathcal{R} with $p, q \in \Pi$, and $\{t_0\}$ an initial state of \mathcal{R} . Let $\tilde{\mathcal{R}}, \tilde{L}_\Pi$, and $\{\tilde{t}_0\}$ be the result of the BR-transformation applied to \mathcal{R}, L_Π , and $\{t_0\}$. Then*

$$\mathcal{R}, L_\Pi, \{t_0\} \models \Box (p \rightarrow (\Diamond_{\leq r} q)) \quad \longrightarrow \quad \tilde{\mathcal{R}}, \tilde{L}_\Pi, \{\tilde{t}_0\} \models \Box (\text{clock}(c_{BR}) \leq r),$$

and

$$\tilde{\mathcal{R}}, \tilde{L}_\Pi, \{\tilde{t}_0\} \models (\Box (p \rightarrow (\Diamond q))) \wedge (\Box (\text{clock}(c_{BR}) \leq r)) \quad \longrightarrow \quad \mathcal{R}, L_\Pi, \{t_0\} \models \Box (p \rightarrow (\Diamond_{\leq r} q)),$$

where $\text{clock}(c_{BR})$ denotes the value of the clock attribute of the clock object c_{BR} .

Proof. For the first statement we show that

$$\tilde{\mathcal{R}}, \tilde{L}_\Pi, \{\tilde{t}_0\} \not\models \Box (\text{clock}(c_{BR}) \leq r)$$

implies

$$\mathcal{R}, L_\Pi, \{t_0\} \not\models \Box (p \rightarrow (\Diamond_{\leq r} q)).$$

Thus assume $\tilde{\mathcal{R}}, \tilde{L}_\Pi, \{\tilde{t}_0\} \not\models \Box (\text{clock}(c_{BR}) \leq r)$. That means, there exists a path $\tilde{\pi} = \{\tilde{t}_0\} \xrightarrow{r_0} \{\tilde{t}_1\} \xrightarrow{r_1} \dots$ of $\tilde{\mathcal{R}}$ with $\tilde{t}_i = t_i < c_{BR} : \text{Clock} \mid \text{clock} : x_i, \text{status} : y_i >$ and a smallest index j such that $x_j > r$. Since the clock value is initially 0 and it increases only due to tick rules if the clock is on, the clock must have been switched on at some point before j . Furthermore, since j is minimal, the clock is continuously on from the last point where it was switched on till \tilde{t}_j .

Assume $i < j$ to be the smallest index such that the clock is continuously on from \tilde{t}_i till \tilde{t}_j . Either i is 0 and the initial state satisfies $p \wedge \neg q$ and $x_i = 0$, or $i > 0$ and the rewrite from the $(i-1)$ th state to the i th

state switched the clock from off to on and reset it to 0. In the latter case the corresponding rewrite has the condition that $p \wedge \neg q$ holds in the i th state. Thus $p \wedge \neg q \wedge x_i = 0$ holds in state \tilde{t}_i . The clock was kept on from state \tilde{t}_i till state \tilde{t}_j . The only rules yielding this behavior are the tick rules increasing the clock value with the duration of the rewrite, and instantaneous rules assuring the invariance of $\neg q$ and letting the clock value untouched. Due to tick-invariance, tick rules cannot cause any change in the validity of the propositions, and $\neg q$ holds all the way from the i th till the j th state. Furthermore, the clock value at state j is the sum of the durations of the rewrites from the i th to the j th state. Thus

$$\exists i, j. 0 \leq i < j \wedge (\tilde{\pi}^i \models p) \wedge \left(\forall i \leq k \leq j. \tilde{\pi}^k \not\models q \right) \wedge \sum_{k=i}^{j-1} r_k > r$$

holds and with Lemma 4.2 we get $\tilde{\pi} \not\models \square (p \rightarrow (\diamond_{\leq r} q))$. Using Lemma 4.1 we conclude that there is also a path π of \mathcal{R} such that $\pi \not\models \square (p \rightarrow (\diamond_{\leq r} q))$ and thus $\mathcal{R}, L_{\Pi}, \{t_0\} \not\models \square (p \rightarrow (\diamond_{\leq r} q))$.

For the second statement assume that

$$\mathcal{R}, L_{\Pi}, \{t_0\} \not\models \square (p \rightarrow (\diamond_{\leq r} q))$$

holds. We show that it implies

$$\tilde{\mathcal{R}}, \tilde{L}_{\Pi}, \{\tilde{t}_0\} \not\models (\square(p \rightarrow (\diamond q))) \wedge (\square(\text{clock}(c_{BR}) \leq r)) .$$

Due to the assumption there exists a path $\pi = \{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots$ of \mathcal{R} violating $\square (p \rightarrow (\diamond_{\leq r} q))$. Now, either $\tilde{\mathcal{R}}, \tilde{L}_{\Pi}, \{\tilde{t}_0\} \not\models \square(p \rightarrow (\diamond q))$ and we are ready, or due to Lemma 4.1 there exists a path $\tilde{\pi} = \{\tilde{t}_0\} \xrightarrow{r_0} \{\tilde{t}_1\} \xrightarrow{r_1} \dots$ of $\tilde{\mathcal{R}}$ also violating $\square (p \rightarrow (\diamond_{\leq r} q))$. With Lemma 4.2 we get

$$\exists i, j. 0 \leq i < j \wedge (\tilde{\pi}^i \models p) \wedge \left(\forall i \leq k \leq j. \tilde{\pi}^k \not\models q \right) \wedge \sum_{k=i}^{j-1} r_k > r.$$

Let i and j be the smallest indices satisfying the above condition.

- If $i = 0$ then by the fact that $\tilde{\pi}^i \models p \wedge \neg q$ we have by definition that the clock in \tilde{t}_0 is on and has the value 0.
- If $i > 0$ and for all $n < i$, \tilde{t}_n does not satisfy $p \wedge \neg q$, then by definition of the initial state the clock is initially off and the clock does not get switched on until the $(i-1)$ th state, thus the clock is off in the $(i-1)$ th state.
- If $i > 0$ and there is an $n < i$ with \tilde{t}_n satisfying $p \wedge \neg q$, then from the minimality of i we conclude that there is a minimal $n \leq m < i$ such that \tilde{t}_m satisfies q . From the minimality of m we conclude that $\{\tilde{t}_{m-1}\} \xrightarrow{r_{m-1}} \{\tilde{t}_m\}$ is due to an instantaneous rule, which, by definition, switches the clock off.

Thus either $i = 0$ and the clock is on in \tilde{t}_i with value 0, or $i > 0$ and the clock is off in state \tilde{t}_{i-1} . Furthermore, in the latter case the $(i-1)$ th state satisfies $\neg p \vee q$ (otherwise i would not be minimal), and the rewrite $\{\tilde{t}_{i-1}\} \xrightarrow{r_i} \{\tilde{t}_i\}$ is due to an instantaneous rule, which, again by definition, switches the clock on and resets its value to 0.

We get that the clock is on with value 0 in \tilde{t}_i . As $\neg q$ holds all the way from the i th till the j th state, the clock remains on from the i th till the j th state. The rewrites of $\tilde{\mathcal{R}}$ assure that the clock value in state \tilde{t}_j is the duration $\sum_{k=i}^{j-1} r_k$ that is by assumption larger than r , what was to be shown. ■

The following lemma states that finiteness of the state space is preserved under the *BR*-transformation, implying that our bounded response model checking algorithm terminates for finite-space systems.

Lemma 4.4. *Given a real-time rewrite theory \mathcal{R} , a labeling function L_Π of \mathcal{R} with $p, q \in \Pi$, an initial state $\{t_0\}$ of \mathcal{R} , and a fixed time sampling strategy, and furthermore, assuming that*

- *there are only finitely many states reachable in \mathcal{R} from initial state $\{t_0\}$ with the given time sampling, i.e., the set*

$$\{\{t_i\} \mid \pi = \{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots \in \text{Paths}(\mathcal{R})_{t_0}, i \in \mathbb{N}\}$$

is finite, and

- *the number of different rewrite durations in all possible paths in \mathcal{R} from $\{t_0\}$ under the given time sampling is finite, i.e., the set*

$$\{r_i \mid \pi = \{t_0\} \xrightarrow{r_0} \{t_1\} \xrightarrow{r_1} \dots \in \text{Paths}(\mathcal{R})_{t_0}, i \in \mathbb{N}\}$$

is finite,

then the bounded response model checking algorithm for \mathcal{R} using the same sampling strategy terminates.

Proof. Assume that the above conditions hold. Notice that the bounded response model checking algorithm always terminates if the set of reachable states of the *BR*-transformation (from its initial state and under the given time sampling) is finite.

Since all instantaneous rules in the *BR*-transformation $\tilde{\mathcal{R}}$ either leave the clock value untouched or reset the clock value to 0, the finiteness of the state space is preserved under the instantaneous rules of $\tilde{\mathcal{R}}$. For the tick rules, on the one hand, if the clock value gets larger than the bound r in the bounded response formula, then the model checking algorithm finds a counterexample and thus terminates. On the other hand, since there are only finitely many possible rewrite durations, there are only finitely many possible clock values less than or equal to r . So if the clock value never exceeds r then the reachable state space of the *BR*-transformation remains finite and the algorithm terminates in this case, too. ■

5 Case Studies

This section briefly presents two case studies where we use the new model checking commands. The analysis has been performed on a 2.4GHz Intel® Core 2 Duo processor with 2 GB of RAM.

5.1 A Network of Medical Devices

We apply the new Real-Time Maude commands on a Real-Time Maude model of an interlock protocol for a small network of medical devices, integrating an X-ray machine, a ventilator machine, and a controller. The example was proposed by Lui Sha, and the Real-Time Maude model is explained in [14].

The ventilator machine helps a sedated patient to breathe during a surgery. An X-ray can be taken during the surgery by pushing a button. To allow an X-ray to be taken without blurring the picture, the ventilator must be briefly turned off. Within a certain time bound, the X-ray must be taken and then the ventilation machine must be restarted. Furthermore, the ventilation machine should not be stopped too often. The model also addresses nondeterministic message delays and clock *drifts*.

In this model, all events take place when some “timer” expires or when a message arrives. Therefore, as proved in [19], the system can be analyzed using the *maximal* time sampling strategy which advances time until the next timer expires, so that the analyses remain sound and complete. One time unit in the specification corresponds to one millisecond in the case study.

Bounded Response Analysis. One requirement in this model is that “the ventilation machine should not pause for more than two seconds at a time.” This can be expressed by the bounded response formula

$$\Box (\text{“machine is pausing”} \longrightarrow \Diamond_{\leq 2\text{sec}} \text{“machine is breathing”}).$$

In order to analyze this property, we first define two state propositions, `isPausing` and `isBreathing`, in the expected way: `isPausing` holds for states in which the ventilation machine is not breathing, while `isBreathing` holds when the ventilation machine is breathing. The bounded response property is model checked using the following Real-Time Maude command:

```
Maude> (br initState /= isPausing => <>le( 2000 ) isBreathing .)
```

The result of this command is a path representing a counterexample to the validity of the property:

Property not satisfied

Counterexample path:

```
{< ct : Controller | clock : 0, lastPauseTime : 0 >
< u : User | pushButtonTimer : 0, pushInterval : 60000 >
< vm : VentMachine | state : breathing >
< xr : X-ray | state : idle >}

=>[pushButton]

{< ct : Controller | clock : 0, lastPauseTime : 0 >
< u : User | pushButtonTimer : 60000, pushInterval : 60000 >
< vm : VentMachine | state : breathing >
< xr : X-ray | state : idle >
dly(pushButton,0,50,10)}

=>[dlyMsgArrives]

...

=>[idle]

{< ct : Controller | clock : 44000/21, lastPauseTime : 3000 >
< u : User | pushButtonTimer : 1220000/21, pushInterval : 60000 >
< vm : VentMachine | state : stopBreathing(9000/7)>
< xr : X-ray | state : idle >}

=>[tick]

{< ct : Controller | clock : 11000/3, lastPauseTime : 3000 >
< u : User | pushButtonTimer : 170000/3, pushInterval : 60000 >
< vm : VentMachine | state : stopBreathing(0)>
< xr : X-ray | state : idle >}
```

The result shows that the bounded response requirement does not hold. This is due to the fact that the ventilation machine may pause for 2.22 seconds, since its internal clock is a little slow (see [14]). A counterexample path is therefore produced, of which we display here only a part, showing the sequence of rules that have been applied to reach a state where the clock added internally to the system reaches a clock value greater than 2000. The analysis took less than a second to perform.

A similar analysis can be done to check whether the ventilation machine cannot pause for more than 2.5 seconds. Since this property holds, the execution of the bounded response command will simply not stop, since the state space reachable from the initial state is not finite (i.e. due to the controller clock attribute, which just increases as time advances).

Minimum Separation Analysis. Another requirement says that the ventilator cannot pause more than once in ten minutes. That is, the *minimum separation* between two pauses is ten minutes. This property can be model checked in Real-Time Maude as follows:

```
Maude> (ms initState /= isPausing separated by >= 600000 .)
```

Property not satisfied

Counterexample path:

```
{< ct : Controller | clock : 0, lastPauseTime : 0 >
< u : User | pushButtonTimer : 0, pushInterval : 60000 >
< vm : VentMachine | state : breathing >
< xr : X-ray | state : idle >}

=>[pushButton]

...

=>[stopBreathing]

{< ct : Controller | clock : 5951000/9, lastPauseTime : 663000 >
< u : User | pushButtonTimer : 530000/9, pushInterval : 60000 >
< vm : VentMachine | state : stopBreathing(2000)>
< xr : X-ray | state : wait(2500/3)>}
```

The requirement does not hold and a counterexample path is produced in less than 10 secs, leading to a state where the internal Clock object reaches a clock value smaller than 600000, while its status is off.

5.2 A Four-Way Traffic Intersection System

In this section, we analyze a bounded response property of an object-oriented Real-Time Maude model of a distributed fault-tolerant four-way traffic light controller for cars and pedestrians described in [17]. The traffic light system for the 4-way intersection is designed as a collection of autonomous concurrent objects that interact with each other by asynchronous message passing. The system is highly parametric: ten different parameters can be specified for an initial state, such as the presence of failures or emergency vehicles in the environment. Each 4-way intersection has two roads crossing in two directions: east-west (EW in the specification) and north-south (NS in the specification). Each road has its own traffic lights. Each pedestrian light has a button that can be pushed by a pedestrian in order to get the green light and cross the street. The behavior of the four-way intersection is as expected.

We focus on the requirement that “no pedestrian should wait for more than five minutes” to cross a road. This corresponds to the bounded response formula

$$\Box (\text{“pedestrian pushes the button”} \longrightarrow \Diamond_{\leq 5min} \text{“pedestrian light is green”}).$$

In order to analyze this property, we use the state propositions `buttonPushed` and `pedLightGreen` that take as parameter the direction of the crosswalk. In less than 3 minutes, we successfully verified that

the pedestrian does not have to wait for more than 15 time units by executing the following Real-Time Maude command (a time unit corresponds to 15 seconds):

```
Maude > (br init("Imoan", minGreenTime + 2, minRedTime, 0, 0, 0, 1, 1, false, 0)
        /= buttonPushed(NS) => <>le( 15 ) pedLightGreen(NS) .)
```

Property satisfied

Furthermore, executing the same command, but for 14 time units, returned a counterexample.

6 Related Work

There are several works determining decidable fragments of timed temporal logics (e.g., [7, 23]) in order to support model checking algorithms for real-time systems. The tools KRONOS [27] and REDLIB [26] are two TCTL (timed CTL) model checkers for timed automata. The popular timed-automaton-based tool UPPAAL [5] provides model checking only for a “reachability subset” of TCTL that does not include bounded response or minimum separation.

The contrast to our work is already explained in the introduction. Whereas the timed automaton formalism is quite restrictive for the exact purpose of achieving decidability of analyses, Real-Time Maude, and even its flat object-oriented subset considered in this paper, is a much more expressive model. The cost of this expressiveness is of course that most properties are in general undecidable for Real-Time Maude. So also for the model checking commands in this paper, which are not guaranteed to terminate for many Real-Time Maude models. Furthermore, since for dense time, Real-Time Maude executes the tick rules according to a time sampling strategy, we must also prove that, even when terminating, our model checking analyses are both sound and complete, using, e.g., the techniques in [19]. Another obvious difference is that we are covering only a fairly small, but important, subset of a MTL.

7 Concluding Remarks

This paper has explained how we have enriched the important class of flat object-oriented Real-Time Maude models with model checking features for bounded response and minimum separation properties.

Object-oriented Real-Time Maude specifications capture many systems that cannot be specified as timed automata; indeed, all advanced Real-Time Maude applications have been so specified. It is therefore not surprising that the model checking problems we address are undecidable in general. Therefore, our model checking analyses may fail to terminate, although they will terminate if the properties do *not* hold. Furthermore, our model checking commands are executed with a selected time sampling strategy, so that only a subset of all possible behaviors are analyzed. Hence, our analyses may be incomplete or unsound. Nevertheless, for object-oriented specifications we have identified easily checkable conditions that ensure soundness and completeness of (untimed) model checking. Further on the positive side, we have shown that (with reasonable assumptions on the treatment of dense time), our model checking analyses terminate when the reachable state space is finite.

The implementation of our model checking procedures follows a transformational approach that takes advantage of Maude’s high performance search command by transforming an MTL model checking problem into checking the validity of an invariant property. We proved the correctness of these transformations under mild conditions, such as tick-invariance and time divergence.

The model checking commands have been integrated into Real-Time Maude and have been successfully used to model check a small network of medical devices [14], as well as on a larger model of a traffic intersection system [17].

The present work is just our first foray into model checking metric temporal logic properties for Real-Time Maude specifications. Much work remains ahead. First of all, we should extend the class of MTL formulas we can model check, and extend the classes of Real-Time Maude models for which such model checking can be performed. For example, if the present techniques could be extended to *non-flat* (or *hierarchical* “Russian dolls”) object-oriented Real-Time Maude specifications, then we would get for free model checkers for these properties for both behavioral AADL models and hierarchical Ptolemy II DE models. We should also extend the commands to analyze only paths up to a certain duration, so that the reachable state space becomes finite. The correctness proofs in this paper all deal with correctness w.r.t. the executed paths. We must of course further investigate the soundness and completeness of such analyses w.r.t. all possible behaviors of a system.

Acknowledgments. We thank the anonymous reviewers for very helpful comments on a previous version of this paper, and gratefully acknowledge financial support by the Research Council of Norway through the Rhythm project, and by the Research Council of Norway and the German Academic Exchange Service (DAAD) through the DAADppp project “Hybrid Systems Modeling and Analysis with Rewriting Techniques (HySmart).”

References

- [1] M. AlTurki, D. Dhurjati, D. Yu, A. Chander & H. Inamura (2009): *Formal Specification and Analysis of Timing Properties in Software Systems*. In: *Proc. of the 12th Int. Conf. on Fundamental Approaches to Software Engineering (FASE’09)*, LNCS 5503, Springer-Verlag, pp. 262–277.
- [2] M. AlTurki & J. Meseguer (2007): *Real-Time Rewriting Semantics of Orc*. In: *Proc. of the 9th ACM SIGPLAN Int. Conf. on Principles and Practice of Declarative Programming (PPDP’07)*, ACM, pp. 131–142.
- [3] R. Alur & T.A. Henzinger (1992): *Logics and Models of Real Time: A survey*. In: *Real Time: Theory in Practice*, LNCS 600, Springer-Verlag, pp. 74–106.
- [4] K. Bae, P. C. Ölveczky, T. H. Feng & S. Tripakis (2009): *Verifying Ptolemy II Discrete-Event Models Using Real-Time Maude*. In: *Proc. of the 11th Int. Conf. on Formal Engineering Methods (ICFEM’09)*, LNCS 5885, Springer-Verlag, pp. 717–736.
- [5] G. Behrmann, A. David & K. G. Larsen (2004): *A Tutorial on UPPAAL*. In: *Proc. of the 4th Int. School on Formal Methods for the Design of Computer, Communication and Software Systems: Real Time (SFM-RT’04)*, LNCS 3185, Springer-Verlag, pp. 200–236.
- [6] A. Boronat & P. C. Ölveczky (2010): *Formal Real-Time Model Transformations in MOMENT2*. In: *Proc. of the 13th Int. Conf. on Fundamental Approaches to Software Engineering (FASE’10)*, LNCS, Springer-Verlag. To appear.
- [7] P. Bouyer (2009): *From Qualitative to Quantitative Analysis of Timed Systems*. Ph.D. thesis, Université Paris.
- [8] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer & C. Talcott (2007): *All About Maude - A High-Performance Logical Framework*, LNCS 4350. Springer-Verlag.
- [9] H. Ding, C. Zheng, G. Agha & L. Sha (2003): *Automated Verification of the Dependability of Object-Oriented Real-Time Systems*. In: *Proc. of the 9th IEEE Int. Workshop on Object-Oriented Real-Time Dependable Systems (WORDS’03)*, IEEE Computer Society Press, pp. 171–178.

- [10] M. Katelman, J. Meseguer & J. Hou (2008): *Redesign of the LMST Wireless Sensor Protocol through Formal Modeling and Statistical Model Checking*. In: *Proc. of the 10th IFIP Int. Conf. on Formal Methods for Open Object-Based Distributed Systems (FMOODS'08)*, LNCS 5051, Springer-Verlag, pp. 150–169.
- [11] R. Koymans (1990): *Specifying Real-Time Properties with Metric Temporal Logic*. *Real-Time Syst.* 2(4), pp. 255–299.
- [12] D. Lepri, P. Cs. Ölveczky & E. Ábrahám (2010). *Model Checking Classes of Metric LTL Properties of Object-Oriented Real-Time Maude Specifications*. Technical Report. http://www-i2.informatik.rwth-aachen.de/~eab/papers/mtl_checking.pdf.
- [13] E. Lien & P. C. Ölveczky (2009): *Formal Modeling and Analysis of an IETF Multicast Protocol*. In: *Proc. of the 7th IEEE Int. Conf. on Software Engineering and Formal Methods (SEFM'09)*, IEEE Computer Society Press, pp. 273–282.
- [14] P. C. Ölveczky (2008): *Towards Formal Modeling and Analysis of Networks of Embedded Medical Devices in Real-Time Maude*. In: *Proc. of the 9th ACIS Int. Conf. on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD'08)*, IEEE Computer Society Press, pp. 241–248.
- [15] P. C. Ölveczky, A. Boronat & J. Meseguer (2010): *Formal Semantics and Analysis of Behavioral AADL Models in Real-Time Maude*. In: *Proc. FMOODS/FORTE'10*. To appear.
- [16] P. C. Ölveczky & M. Caccamo (2006): *Formal Simulation and Analysis of the CASH Scheduling Algorithm in Real-Time Maude*. In: *Proc. of the 9th Int. Conf. on Fundamental Approaches to Software Engineering (FASE'06)*, LNCS 3922, Springer-Verlag, pp. 357–372.
- [17] P. C. Ölveczky & J. Meseguer: *Specification and Verification of Distributed Embedded Systems: A Traffic Intersection Product Family*. To appear in *Proc. RTRTS 2010*.
- [18] P. C. Ölveczky & J. Meseguer (2002): *Specification of Real-Time and Hybrid Systems in Rewriting Logic*. *Theoretical Computer Science* 285, pp. 359–405.
- [19] P. C. Ölveczky & J. Meseguer (2007): *Abstraction and Completeness for Real-Time Maude*. *Electronic Notes in Theoretical Computer Science* 176(4), pp. 5–27.
- [20] P. C. Ölveczky & J. Meseguer (2007): *Semantics and Pragmatics of Real-Time Maude*. *Higher-Order and Symbolic Computation* 20(1-2), pp. 161–196.
- [21] P. C. Ölveczky, J. Meseguer & C. L. Talcott (2006): *Specification and Analysis of the AER/NCA Active Network Protocol Suite in Real-Time Maude*. *Formal Methods in System Design* 29(3), pp. 253–293.
- [22] P. C. Ölveczky & S. Thorvaldsen (2009): *Formal Modeling, Performance Estimation, and Model Checking of Wireless Sensor Network Algorithms in Real-Time Maude*. *Theoretical Computer Science* 410(2-3), pp. 254–280.
- [23] J. Ouaknine & J. Worrell (2005): *On the Decidability of Metric Temporal Logic*. In: *Proc. of the 20th Annual IEEE Symposium on Logic in Computer Science (LICS '05)*, IEEE Computer Society Press, pp. 188–197.
- [24] A. Pnueli (1977): *The Temporal Logic of Programs*. In: *Proc. of the 18th Annual Symposium on Foundations of Computer Science (SFCS'77)*, IEEE Computer Society Press, pp. 46–57.
- [25] J. E. Rivera, F. Durán & A. Vallecillo (2010): *On the Behavioral Semantics of Real-Time Domain Specific Visual Languages*. In: *Proc. of the 8th Int. Workshop on Rewriting Logic and its Applications (WRLA'10)*, LNCS, Springer-Verlag. To appear.
- [26] F. Wang (2006): *REDLIB for the Formal Verification of Embedded Systems*. In: *Proc. of the 2nd Int. Symposium on Leveraging Applications of Formal Methods (ISoLA'06)*, IEEE Computer Society Press, pp. 341–346.
- [27] S. Yovine (1997): *Kronos: A Verification Tool for Real-Time Systems*. *Software Tools for Technology Transfer* 1(1-2), pp. 123–133.